Autofolding for Source Code Summarization

Jaroslav Fowkes[†] Razvan Ranca[‡] Miltiadis Allamanis[†]
Mirella Lapata[†] Charles Sutton[†]

[†]School of Informatics
University of Edinburgh, Edinburgh, EH8 9AB, UK
{jfowkes, m.allamanis, mlap, csutton}@ed.ac.uk

[‡]Computer Laboratory
University of Cambridge, Cambridge, CB3 0FD, UK
rr463@cam.ac.uk

arXiv:1403.4503v1 [cs.SE] 18 Mar 2014

ABSTRACT

We present a novel fully unsupervised method for the automatic summarization of source code text. On large projects, developers spend much of their time reading and browsing code, raising new opportunities for summarization methods. Indeed, modern code editors contain a feature called *code folding*, which allows one to selectively hide blocks of code. But this feature is cumbersome because currently folding decisions must be done manually.

We introduce the *autofolding* problem, which is to automatically create a code summary by folding non-essential code blocks. We present a novel solution to the autofolding problem by formulating it as a subtree optimization problem and making use of a comprehensive source code content model. On a set of popular open source projects from GitHub, we show that our formulation significantly outperforms simpler baselines at meeting human judgements, yielding an 84% accuracy and a 43% error reduction. More broadly, we hope that this work will aid program comprehension by turning code folding into a usable and valuable tool.

1. INTRODUCTION

Engineering large software systems presents many challenges due to the inherent complexity of software. Because of this complexity, programmers tend to spend more time reading and browsing code than actually writing it [22, 20]. Despite much research [42], there is still a large need for better tools that aid program comprehension, thereby reducing the cost of software development.

A key insight is that source code is written to be understood not only by machines, but also by humans. Programmers devote significant time and attention to writing their code in an idiomatic and intuitive way that can be easily understood by others — *source code is a means of human communication.* This fact raises the intriguing possibility that technology from the natural language processing (NLP) community can be adapted to help developers make sense of large repositories of code. A common problem in program comprehension is that developers are faced with two extremes: skimming the code or reading it in depth [14]. A good *summary* of the source code would alleviate this problem by providing a compromise: it could be read quickly and would lead to better understanding than simply skimming the code.

Source code summarization has potential for valuable applications in many software engineering tasks, such as: (a) *Understanding new code bases.* Often developers need to quickly familiarize themselves with the core parts of a large code base. This can happen when a developer is joining an existing open source project, or when a developer is evaluating whether to use a new software library. (b) *Code reviews.* Reviewers need to quickly understand the key changes before reviewing the details. (c) *Locating relevant code segments.* During program maintenance, developers often tend to skim source code, reading only a couple lines at a time, while searching for the code region they are interested in [40].

For this reason, many code editors include a feature called *code folding*, which allows developers to selectively display or hide blocks of the source code. This feature is commonly supported in editors and is familiar to developers [16, 21, 35]. But in current Integrated Development Environments (IDEs), folding quickly becomes cumbersome because the folding decisions must be done manually by the programmer. This creates an obvious chicken-and-egg problem, because the developer must already understand the source file to decide what should be folded.

In this paper, we propose that code folding can be a valuable tool for aiding program comprehension, provided that folding decisions are made automatically. We introduce the *autofolding problem*, in which the goal is to automatically create a code summary by folding non-essential code elements that are not useful on first viewing. An illustrative example is shown in Figure 1. To any Java developer the function of the `StatusLine` constructor and the `clone`, `getCode`, `getReason` and `toString` methods are obvious even without seeing their method bodies. One possible summary of this source file is shown in Figure 2.

Our contributions in this paper are as follows:

- We introduce a novel autofolding method for source code summarization based on optimizing the similarity between the summary and the full source file. Because of certain constraints among the folding decisions, we formulate this method as a contiguous rooted subtree problem (Section 3.3). This is, to our knowledge, the first method for code summarization to use autofolding.
- In order to determine which non-essential code regions our summarizer should fold, we adapt a content model from the NLP community and thereby create a novel topic model for code that is of independent interest (Section 3.2). The key feature of our topic model is that it is scoped, allowing every project and file to be endowed with its own particular topic.
- We perform a comprehensive evaluation of our method on a set of popular open source projects from GitHub

```java
/* Header */
package org.zoolu.sip.header;

/** SIP Status-line, i.e. the first
 * line of a response message */
public class StatusLine {
 protected int code;
 protected String reason;

 /** Construct StatusLine */
 public StatusLine(int c, String r) {
  code = c;
  reason = r;
 }

 /** Create a new copy of the request-line */
 public Object clone() {
  return new StatusLine(getCode(), getReason());
 }

 /** Indicates whether some other Object
  * is "equal to" this StatusLine */
 public boolean equals(Object obj){
  try {
   StatusLine r = (StatusLine) obj;
   if (r.getCode() == getCode()&&
   r.getReason().equals(getReason()))
    return true;
   else
    return false;
  } catch (Exception e) {
   return false;
  }
 }

 public int getCode() {
  return code;
 }

 public String getReason() {
  return reason;
 }

 public String toString() {
  return "SIP/2.0 " + code + " " + reason + "\r\n";
 }
}
```

Figure 1: Original source code. A snippet from `bigbluebutton`'s `StatusLine.java`. We use this as a running example of code summarization.

(Section 4), and find that our formulation performs significantly better than simpler baselines (Section 5). Furthermore, our proposed method is able to match human judgements with an accuracy of 84%.

More broadly, we hope that this work will aid program comprehension by turning code folding, perhaps an overlooked feature, into a useful, usable and valuable tool.

2. RELATED WORK

The application of NLP methods to the analysis of source code text is only just beginning to be explored. Recent work has applied language modelling [2, 17, 18, 23, 32], natural language generation [38, 39], machine translation [31], and topic modelling [30] to the text of source code from large software projects.

One of the main challenges in this area is to adapt existing NLP techniques to source code text. In contrast to natural languages, programming languages are unambiguous, employ little redundancy, are meant to be interpreted literally, and consist of strictly structured text. To exploit these features of the problem, we perform the summarization at the code block level, leveraging the fact that source code is syntactically unambiguous.

```java
/* Header...*/
package org.zoolu.sip.header;

/** SIP Status-line, i.e. the first...*/
public class StatusLine {
 protected int code;
 protected String reason;

 /** Construct StatusLine...*/
 public StatusLine(int c, String r) {...}

 /** Create a new copy of the request-line ..*/
 public Object clone() {...}

 /** Indicates whether some other Object...*/
 public boolean equals(Object obj){
  try {
   StatusLine r = (StatusLine) obj;
   if (r.getCode() == (getCode()&&
   r.getReason().equals(getReason()))
    return true;
   else
    return false;
  } catch (Exception e) {...}
 }

 public int getCode() {...}

 public String getReason() {...}

 public String toString() {...}
}
```

Figure 2: Summarized source code. A summary of the file in Figure 1 (left) which results from folding lines 1, 4–5, 11–14, 21–22, 31–33, 36-38 and 40-42.

There is some existing work on the use of manual code folding to aid comprehension. In particular, Rugaber et al. [35] consider a conceptual model for manual folding, extending it to non-contiguous regions of code. Kullbach et al. [21] develop the `GUPRO` IDE to aid in the comprehension of C preprocessor code by folding macro expansions and file includes. Also, Hendrix et al. develop the `GRASP` program comprehension tool, combining control structure diagramming with folding [16].

However, to best of our knowledge, the autofolding problem is novel. We are aware of only a few previous methods that consider the problem of summarizing code artifacts. Most similar to our work are Haiduc et al. [14, 15] and the follow up work by Eddy et al. [9], who also consider the problem of summarizing source code, particularly methods and classes, but in their work code fragments are summarized by a short list of keywords. For example, the `equals` method in Figure 1 might be summarized by the list of terms (`equals`, `code`, `reason`, `Status`). In our work, we summarize code *with* code, which we would argue has the potential to provide a much richer and more informative summary.

Also, Ying et al. [47] consider the problem of summarizing a list of code fragments, such as those returned by a code

search engine. They use a supervised learning approach at the level of lines of code. Because they consider the results of code search, their classifier uses query-level features, e.g., whether a line of code uses any identifiers that were present in the query. This is a source of information that is not available in our problem setting. In contrast, we target use cases in which the developer is skimming the source code to get an overview of its operation, rather than performing a directed keyword search. Kim et al. [19] develop a system that augments API documentation with code example summaries but these are mined from the web and are therefore limited to APIs which have examples already written for them — our summarization approach is applicable to *any* source file.

On a more technical level, our work is distinguished from this previous work in our use of content selection at the level of syntactic blocks. This results in a summary that is coherent with respect to the programming language's syntax. The previous approaches perform content selection at the line level or the term level, running the risk of producing summaries that consist of disjoint sections of code. Indeed, Eddy et al. [9] observe that developers prefer summaries with a natural structure. Folding on code blocks also enables us to retain method headers in the summary — identified by Haiduc et al. [15] as highly relevant to developers and accounting for the high scores of their best performing method.

Additionally, our method leverages a cross-project corpus during the summarization process, allowing us to detect coding patterns which are common among many projects and which are therefore less likely to be part of the core logic of any one project.

The task of natural language summarization has been studied extensively [37], mostly focusing on *extractive* summarization — the problem of extracting the most relevant text segments from documents. Source code identifiers (e.g., variable names) are information-rich and have been shown to be important for tasks such as feature location [1, 7]. NLP techniques have been used on these identifiers, for information retrieval tasks such as automatically selecting labels for software artifacts [5]. Extractive summarization has also been applied for the automatic summarization of bug reports [25, 34].

In addition to extractive summarization methods, *abstractive* techniques have also been used in software engineering research. Work in this area includes the synthesis of API usage examples [4], the extraction of API usage patterns [45, 46], and the generation of natural language summaries for source code [29, 38].

The use of topic models for source code has also been studied in depth [5, 11, 36, 43]. Marcus et al. [26, 27] used Latent Semantic Indexing (LSI)[6] for identifying traceability links and concept location in source code. More closely related to our work, Haiduc et al. [15] used LSI as a content model for their keyword-based source code summarizer. In their follow up paper, Eddy et al. additionally used a hierarchical pachinko allocation model (hPAM)[28], a family of generative topic models that build on Latent Dirichlet Allocation (LDA)[3] with a hierarchical topic structure. Also, Movshovitz et al. [30] successfully used LDA and link-LDA[10] for predicting class comments from source file text.

We cast autofolding as an instance of the general problem of selecting an optimal subtree given a certain budget. This problem has been studied theoretically by [12], who propose a dynamic programming solution, but this is only pseudo-polynomial time, and so is unlikely to scale well in practice.

3. PROBLEM FORMULATION

Our aim in this paper is to summarize source code so that it conveys the most important aspects of its intended function. We envisage our proposed Tree-based Software Summarization Algorithm (TSSA) being embedded in a programming language Integrated Development Environment (IDE), providing real-time summaries to the user of selected files. The summarization could be useful at multiple levels, ranging from a single source file to a whole corpus. For the purposes of this paper, we will focus on the Java programming language as it is a popular, high-level, platform-independent language. However, since TSSA works entirely with the source code's Abstract Syntax Tree (AST), it can be applied to *any* programming language.

The usage scenario we envisage for TSSA is that of a developer not familiar with a project wishing to see the most relevant methods and classes in a given project file. For example, a developer who is considering using a new project on GitHub and would like to get an overview of the algorithms used in the project. In the future, we would like to extend this to aid developers undertaking other software engineering tasks such as program maintenance, exception handling and code reviews.

3.1 Problem Definition

Most modern IDEs and text editors already have extensive support for *manual* folding of code but to the best of our knowledge the problem of doing so *automatically* is novel. When we say that we *fold* a source code region we mean that the region is replaced by a one line summary and a symbol indicating that the region was folded. We define the *autofolding problem* as that of choosing a set of code regions to fold, such that the total length of the folded file as a fraction of the original is below a user-specified compression ratio, and the remaining, unfolded, text captures the most important aspects of the file in question.

To provide the most intuitive summary to the end-user, we allow the system to preform folding only on *code blocks* (regions of source code delimited by `{` , `}`), *comment blocks* (regions delimited by `/*(*)` , `*/`), and `import` statements. We call these the *foldable regions* of the code. Our reasoning for this is that it is a summary many programmers are familiar with as these are precisely the regions that can be manually folded in the majority of modern IDEs. Moreover, code blocks are natural units for extractive summarization since they take advantage of the code structure specified by the programmer. In keeping with the conventions established by manual folding in IDEs, the one line summary we display for a folded region consists of the first non-empty line of the code block, then a symbol denoting the folded region (an ellipsis), and finally the right delimiter of the region. See Figure 2 for an example of how folded text appears in an IDE.

We formalize the autofolding problem by using the AST representation of the source code[1]. Given a program's AST, we define the program's *foldable nodes* as those AST nodes which correspond to a foldable region of code (a code/com-

[1]The AST representation can be built by the Eclipse JDT parser [8].

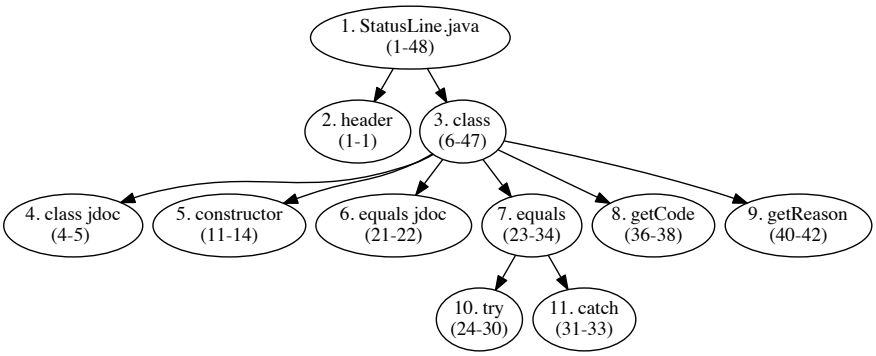

Figure 3: Partial foldable tree constructed for `StatusLine.java` (Figure 1). Numbered breadth-first with labels denoting block types and line numbers in brackets, cf. the source code snippet in Figure 1. Note that we have omitted some nodes for clarity.

ment block or import statement). By starting at the root of an AST and sequentially extracting all foldable nodes, we construct a directed *foldable tree*, containing just the AST nodes we are interested in. Figure 3 shows a partial foldable tree for the running example.

Constructing a foldable tree enables us to formulate the summarization problem mathematically as finding the best contiguous rooted subtree that takes up no more than a predefined number of lines of source code (LOC). That is to say, we unfold all nodes in the best subtree and fold the remaining nodes in the tree. Note that we require the tree to be rooted and contiguous as otherwise this would lead to confusing situations where we would have a deep node present in the summary with no context. We will describe the precise formulation in the next two sections: we begin by presenting a content selection model for choosing the best nodes to retain in the subtree (Section 3.2) before describing a greedy approximation algorithm to solve the contiguous rooted subtree problem (Section 3.3).

3.2 Content Model

In order to determine which nodes of the foldable tree should be unfolded we require a content selection method for choosing the best nodes to retain in the summary. Intuitively, one would like to retain the most informative nodes and a natural approach, as in text summarization, is to tokenize the node text and select the nodes with the most representative tokens. For this reason we make use of an LDA-style *Topic Model*, an extension of the `TopicSum` model [13]. The high-level idea is to extend LDA with topics that are specific to particular projects, files and methods. Note that we do not use an *n-gram language model* as we require a global model of the distribution of code across files rather than a local model based on nearby tokens.

In NLP, a topic model is a generative probabilistic model which captures the distribution of words in documents by associating each word with a number of *topics*. Each topic is modelled as a distribution ϕ_k over words which enables each document to be modelled as a distribution over topics θ_d. Thus, one can generate a document simply by choosing each word w in the document from a topic ϕ_z, where the topic assignment z is selected according to the document's topic distribution θ_d [41].

`TopicSum` is a *scoped* topic model and extends this simple yet powerful idea to multiple levels. The key difference is that each topic can be one of three kinds, a distribution over words in: the document, the document collection and background words. The intention is that the background topic models stop words, the document collection topic represents significant content and the document topic very specific document words [13]. We have adapted `TopicSum` to source code so that each file and project have their own private topic. Further, we have introduced several background topics. Our TSSA content model is best illustrated by example (cf. Figure 4). A file F in our topic model is generated as follows: each token w in every foldable node N of the file is chosen from a topic ϕ_z, where the topic assignment z is selected according to θ_N the distribution over topics in the node. For example, the topic assignment $z = 1$ means that w is chosen from the file token distribution ϕ_F. Specifically, we consider five different kinds of topic: for each file (ϕ_F), for the project containing the file (ϕ_P) and three background topics shared across projects (ϕ_B) as depicted in Figure 4. Although the model isn't aware of this, in practice we find that the three background topics correspond to common Java tokens (ϕ_{B_J}) and background doc comment (ϕ_{B_D}) and header comment tokens (ϕ_{B_H}).

For the purposes of code summarization, we need to evaluate whether a block of code belongs in the summary (cf. Section 3.3) and our means to achieve this goal depends on measuring whether the code is representative of its file as a whole, its project, and of idiomatic source code more generally. The `TopicSum` model is ideally suited to this task and enables us to directly calculate the value of including each node in the summary and assign a 'score' based on this to

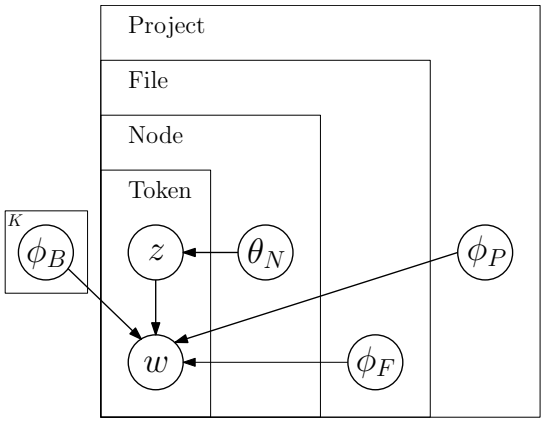

Figure 4: Graphical model depiction of the TSSA content model (with $K = 3$). The plates denote repeated groups of variables. Note that hyperparameters have been omitted for clarity.

each node, as will be described in Section 3.3. Note that we do not use a stoplist because we expect that the set of appropriate stop words for program text would be different to those for natural language text. Moreover, the Java, doc comment and header comment background distributions $\phi_{B_J}, \phi_{B_D}, \phi_{B_H}$ are intended to capture stop words from Java tokens and comment tokens which do not contribute to the content.

We can then construct the topic model by sampling topics from a suitable conditional probability distribution and so infer the topic assignments for our corpus. That is, we consider each individual token in the corpus in turn, and estimate the probability of assigning the current token to each topic, conditioned on the topic assignments to all other tokens. We can then sample a topic from this conditional distribution and store it as the new topic assignment for this token [41]. It can be shown that this process converges to the distribution of topics for the corpus and we are thus able to predict the most likely tokens for each topic.

The conditional distribution used for sampling is a simple extension of the standard equations for LDA [41]: letting $z_i^{(nfp)} = k$ denote the topic assignment of token i (in node n, file f project p) to topic k and $w_i^{(nfp)} = t$ denote token i taking on the value t, we obtain

$$P(z_i^{(nfp)} = k \,|\, \mathcal{Z}_{\setminus i}, \mathcal{W}, \alpha\mathbf{m}, \boldsymbol{\beta}) =$$
$$P(w_i^{(nfp)} = t \,|\, z_i^{(nfp)} = k, \mathcal{Z}_{\setminus i}, \mathcal{W}, \boldsymbol{\beta}) P(z_i^{(nfp)} = k \,|\, \mathcal{Z}_{\setminus i}, \alpha\mathbf{m})$$
$$\propto \frac{N_{t|k}^{\setminus i} + \beta_k}{\sum_{t=1}^{T}\left(N_{t|k}^{\setminus i} + \beta_k\right)} \frac{N_{k|n_{fp}}^{\setminus i} + \alpha m_k}{\sum_{k=1}^{K}\left(N_{k|n_{fp}}^{\setminus i} + \alpha m_k\right)}$$

where $\mathcal{Z}, \mathcal{W}$ denote the assignments of all topics and tokens, $N_{t|k}^{\setminus i}$ denotes the number of times token t is assigned to topic k, not including the current instance i, and similarly $N_{k|n_{fp}}^{\setminus i}$ denotes the number of times topic k is assigned to some token in node n_{fp}, excluding instance i. Further, $\alpha\mathbf{m}, \boldsymbol{\beta}$ denote hyperparameters for the prior distributions of topics and tokens, i.e. the initial topic and token assignments. Additionally, unlike the original `TopicSum` model which used fixed hyperparameters, we incorporate efficient hyperparameter optimization, namely MacKay and Peto's method with histogram based computation of the necessary counts as described in [44].

Tokenization A first idea would be to have one token in the topic model for each Java token in the code base. However there are a few problems with this. First, some tokens, like operators and delimiters, are not informative about the program content. Second, identifier names have substructure that we wish to exploit. For example, in Figure 1, the `getCode` method name is closely related to the `code` member variable, which becomes apparent to the Topic Model only if it is split into two tokens `get` and `code`.

For these reasons, we preprocess the Java tokens before incorporating them into the topic model. Given a code block, first we tokenize it into a set of Java tokens using standard tools for the Java programming language. Then we remove all tokens except for identifiers, i.e., programmer assigned names of variables, methods, classes, etc. Finally, we convert each of the identifiers into a new set of tokens by splitting on camel case and underscores, and converting to lowercase (e.g., `FooBarBaz` becomes three tokens `foo`, `bar`, and `baz`, as would `foo_bar_baz`). Additionally, we include the text of all program comments in the topic model, splitting the comment text based on words, again applying the identifier splitting procedure on any comment tokens that contain camel case or underscores. The resulting tokens are used as the tokens in the vocabulary for the topic model.

3.3 Optimization Method

Now that we have defined a content model for source code, we need an algorithm that extracts the most relevant summary by selecting suitable AST nodes in the foldable tree. To this end, we propose an iterative greedy optimization algorithm to extract the most relevant rooted contiguous subtree from the foldable tree given constraints on the subtree size.

In order to decide whether a given node of the foldable tree is relevant and should be included in our summary, we consider the candidate summary, i.e., the subtree we have gathered so far, together with the node under consideration. Using our topic model, we are then able to assign probabilities to the candidate summary and measure its similarity to the file token distribution ϕ_F. A suitable information-theoretic measure for this is the Kullback-Lieber (KL) divergence, which is a measure of the difference between two probability distributions. We can therefore assign the node a score based on the KL divergence between the corresponding file token distribution ϕ_F and the empirical unigram distribution of the candidate summary. This is similar to the summary selection criterion proposed in the original `TopicSum` model [13], except that Haghighi et al. used the project token distribution and our summary distribution backs off to the Java token distribution ϕ_{B_J} rather than a constant. The intuition here is that we would like our summary to contain nodes with tokens which are representative of the file but not common in the project or corpus as we believe that these tend to be the most relevant.

Formally, for each node i let $u_i \in \{0, 1\}$ indicate whether it is unfolded, with 1 corresponding to true (i.e. unfolded) and let $\mathbf{u}$ be the vector containing all the u_i (so that $\mathbf{u}$ contains u_i at position i). We then define the score as

$$\sigma(\mathbf{u}) = -KL(\phi_F | P_{\mathbf{u}}) \tag{3.1}$$

where ϕ_F denotes the file token distribution, $P_{\mathbf{u}}$ the empirical unigram distribution over tokens in the candidate summary and $KL(\cdot|\cdot)$ the KL divergence. We adjust this score down for nodes consisting of Java *import statements* and *block comments* along with empty blocks so that these are never unfolded as developers generally prefer not to see import statements, file header comments and internal block comments in the summary — the Eclipse IDE for example automatically folds import statements and file header comments by default. We do not, however, impose this condition on *Javadoc comments* as studies [15, 9] have indicated that developers like to see good method/class comments.

We are now in a position to formulate the summarization problem as finding the optimal rooted contiguous subtree. Note that this formulation is necessary as we wish to present a summary of the foldable tree – this requires finding a subtree that is both rooted and contiguous so that its unfolded nodes are all visible in the summary. Suppose we wish to summarize a file using a line-level compression ratio of p%, i.e., we would like to compress the file to p% of its original size in LOC. We can then define L_{max}, the maximum

number of lines of code that are allowed in the summary, as

$$L_{max} := \frac{p}{100} L_0 \tag{3.2}$$

where L_0 denotes the LOC in the original file. Note that since we are folding on a block level, a line-level compression ratio of $p\%$ does not mean that $p\%$ of the blocks are compressed, in fact for small files all the blocks are often folded at 50% compression.

Let $G = (V, E)$ be the foldable tree, that is, a directed tree with a set of AST nodes $V = \{1, \ldots, N\}$ consisting of the aforementioned foldable regions and E a set of *directed* edges between the nodes, where it is understood that $(i, j) \in E$ means that i is the parent node of j (cf. Figure 3). Furthermore, for a node $i \in V$ let L_i denote the LOC underneath node i. We then define C_i, the cost of unfolding node i as the LOC *unique* to node i, i.e., underneath node i but not any of its children. Formally,

$$C_i := (L_i - 1) - \sum_{j:(i,j)\in E} (L_j - 1) \tag{3.3}$$

Recall that the first line of a node is never folded (cf. Figure 2), hence the minus one. Let $\sigma(\mathbf{z})$ denote the score obtained from the summary nodes $u_1, \ldots, u_N$ as defined in (3.1). The optimal rooted contiguous subtree problem is then defined as

$$\max_{\mathbf{u}} \sigma(\mathbf{u}) \tag{3.4a}$$

$$\text{s.t.} \sum_{i=1}^{N} C_i u_i \leq L_{max} \tag{3.4b}$$

$$u_i \geq u_j \qquad \text{if } (i, j) \in E \tag{3.4c}$$

$$u_i \in \{0, 1\} \qquad \forall\, i \in V. \tag{3.4d}$$

That is, we wish to unfold the nodes which maximise the total score (3.4a) subject to staying below the maximum allowed LOC (3.4b) and retaining a rooted contiguous subtree (3.4c).

As the score (3.1) is nonlinear we approximate this problem using a greedy approximation algorithm: that is we unfold the next available node that will give the maximum score (3.1) per cost (3.3) increase, honoring the cost constraint (3.4b). That is, starting from all the nodes being folded (i.e. $\mathbf{u} = \mathbf{0}$), iteratively choose the node that maximizes

$$\sigma(\mathbf{u} + \mathbf{e_i})/C_i$$

while $\sum_{i=1}^{N} C_i u_i \leq L_{max}$. Here $\mathbf{e_i}$ denotes the i-th unit vector, which is 1 at position i and 0 everywhere else.

4. EXPERIMENTAL SETUP

In the previous section, we described source code summarization as a folding problem and presented a greedy algorithm that performs the summarization. In this section, we describe how we obtain a gold standard summary on real code. The gold standard allows automatic evaluation, which is the de facto standard in NLP summarization [24]. Automatic evaluation is vital to applied AI because it leads to a rapid development cycle. In particular, it enables us to perform a comprehensive evaluation of the summarization method in the next section.

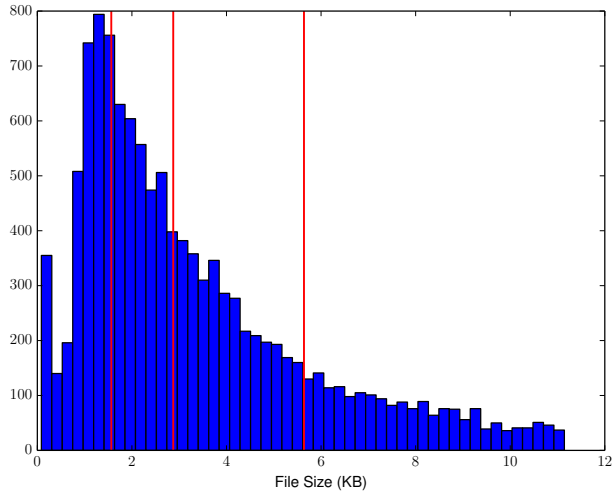

Figure 5: Histogram of file sizes across the top projects (red vertical lines denote quartiles).

Data To evaluate the performance of TSSA we obtained the source code for the top six Java projects on the popular GitHub service for hosting open source code repositories. The top six projects were determined by a popularity score, which is the sum of the number of forks and the number of watchers, where each is separately normalized into a z-score. We selected the projects with the highest score that were greater than 100,000 KB in size as of December 1st 2013. These are given in Table 1 along with a brief description of their domains.

For each of these six projects, we divided the project files into quartiles by file size, and sampled four files from each quartile, to obtain a total of 96 files each one of which was annotated by two human annotators. Only `.java` files were considered (excluding the `package-info.java` files, which only contain package-level documentation). Figure 5 shows the distribution of file sizes across the top projects.

Annotation Procedure Human annotators were given specific guidelines prior to performing the annotation. Annotators were asked, for each source code file, to manually fold the file in the Eclipse IDE until they reached a compression ratio of 75%, save their work, and then to continue folding until they reached a ratio of 50%.[2] Thus each annotator produced two sets of folding decisions for each file. Each annotator annotated all 96 files in our dataset at 75% and 50% compression ratios. Although 50% may not seem like a dramatic compression, in fact many of the remaining lines are block headers or blank (cf. the running example in Figure 2). We found that on average across all 96 files in dataset, over half (56%) of the LOC in the folded file were blank or block headers.

Annotators were allowed to browse the full source code of the project while annotating each file. In order to obtain a high quality gold standard we required human annotators who were conversant with the Java programming language and for this reason we chose two of the authors as annotators. Note that we performed our annotation prior to the development of our summarization system so it was impos-

[2]A compression ratio of 25% was not used as this would result in most files being entirely folded because we are folding on a block level.

Project	Description	LOC	Methods	Classes	Forks	Watchers
storm	Distributed Computation System	59,827	5,740	761	1,416	7,471
elasticsearch	REST search engine	518,905	32,077	4,990	1,283	5,246
spring-framework	Application Framework	798,249	47,214	8,395	1,774	2,568
libgdx	Game Development Framework	334,706	33,821	2,651	1,844	2,243
bigbluebutton	Web Conferencing System	105,315	6,364	852	1,602	969
netty	Network Application Framework	160,579	10,324	1,267	927	2,304

Table 1: The top Java projects on GitHub, used in the current work. Ordered by popularity.

Background			Project		File	
Java	header comments	doc comments	`spring-framework`	`bigbluebutton`	`DataSourceUtils`	`QuaLsp`
`get`	`the`	`the`	`org`	`org`	`connection`	`lsp`
`string`	`license`	`a`	`springframework`	`sip`	`source`	`j`
`name`	`or`	`to`	`bean`	`log`	`data`	`constants`
`value`	`under`	`of`	`test`	`java`	`con`	`k`
`type`	`you`	`is`	`context`	`event`	`holder`	`ld8`
`set`	`distributed`	`link`	`exception`	`message`	`synchronization`	`mode`
`object`	`of`	`for`	`factory`	`gnu`	`transaction`	`index`
`field`	`is`	`this`	`request`	`it`	`isolation`	`tmp`
`to`	`0`	`and`	`java`	`listener`	`order`	`value`
`i`	`2`	`code`	`web`	`public`	`level`	`wegt`

Table 2: The top ten tokens in each topic-type as found by our topic model.

sible for the annotators to unconsciously favor the system's output in their judgements.

Annotators were asked to always fold `import` statements and header comments (such as copyright notices). Empty and one-line blocks were also folded by default. Setters and getters along with other I/O methods were asked to be folded, unless they contained core logic of the code. Similarly, annotators were asked to fold overridden and overloaded methods unless they provide sufficient new information about the functionality of the code. Finally, Javadoc and block comments were left unfolded if they were informative and succinctly explained the function of the associated class or method. Comments whose text spanned only one line were folded by default.

It is certainly true that developers who are familiar with the projects could favor a different type of summary than the annotators, and this could be seen as a threat to validity. However, recall that in this initial study, our target use case focuses on developers who are new to the projects, for which we would argue our annotators are good representatives.

Annotation Statistics Because autofolding is a new task, we need to verify that the task is well-defined. We therefore calculated the agreement between annotators and found that it was substantial, with Fleiss' Kappa values of 0.72 at 75% and 0.64 at 50% compression, averaged across all files.

5. RESULTS

In this section, we evaluate the performance of TSSA against our annotated test set from Section 4. We begin by training and assessing the quality of TSSA's underlying topic model before moving on to a comprehensive evaluation of TSSA itself. As our method is fully unsupervised, we use all 96 annotated files from the six projects in our corpus for evaluation.

Topic Model We train TSSA's topic model on all 12,093 `.java` files from the six projects in our corpus — this enables the topic model to automatically discern common coding patterns, such as common libraries, and thus allows TSSA to recognize unimportant code. To this end, we run the topic sampling algorithm for 5,000 iterations, performing hyperparameter optimization every 10 iterations, to infer our trained model. To highlight the validity of our topic model, we show the top ten tokens in each of the five topic types (Java background, header comment, doc comment, project and file) in Table 2. Note that background tokens which are very common in a file/project can also appear among the top ten file/project tokens. One can clearly see that these are very representative of their respective topics: the Java background topic contains common Java generics such as `get`, `set`, `to` and `object`. The background comment topics on the other hand contain common English stop words such as `the`, `a`, `or` and one can clearly see the distinction between the topics, i.e., words that commonly appear in headers (such as `license`, `distributed`) are found in the header comment topic whereas those that appear in code comments (`link`, `code`) are found in the doc comment topic. Looking at the topics for the `spring-framework` project one can see that it contains the fully resolved project name `org.springframework` as well as common project-specific tokens such as `bean`, `context` and `factory`. Similarly, the `bigbluebutton` project contains the specific tokens `sip`, `event`, `message` and `gnu`. Looking at selected files from both projects, the `DataSourceUtils` file topic contains tokens specific to the function of that file, e.g. `connection`, `data`, `synchronization` and `transaction`. Finally, the `QuaLsp` file topic, a codec implementation, contains the very function-specific tokens `lsp` and `ld8`. This illustrates the quality of our proposed topic model and suggests that it can distinguish file- and project-specific tokens from those that are common across all Java projects. This raises the exciting possibility that our topic model is robust enough for wider applications and not merely restricted to summarization.

Baselines To provide a comprehensive performance comparison of our proposed summarization method, we also evaluate four alternative baseline methods that represent more naïve approaches for summarizing source code. All baselines start from a fully folded tree and gradually unfold nodes, making local decisions, until they reach the required compression ratio. A brief description and example behavior

Method	Accuracy	F_1	Precision	Recall
TSSA	**0.87**	**0.90**	**0.90**	**0.90**
Javadocs	0.71	0.79	0.78	0.81
Shallowest	0.52	0.65	0.64	0.66
Largest	0.54	0.66	0.65	0.67

Table 3: Per-line evaluation statistics for the summarizers at a compression ratio of 75%. Averaged across all test files, ordered best first.

Method	Accuracy	F_1	Precision	Recall
TSSA	**0.84**	**0.83**	**0.85**	**0.81**
Javadocs	0.72	0.69	0.69	0.69
Shallowest	0.60	0.55	0.56	0.53
Largest	0.60	0.54	0.55	0.53

Table 4: Per-line evaluation statistics for the summarizers evaluated at a compression ratio of 50%. Averaged across all test files, ordered best first.

on the folded tree in Figure 3 follows. Note that if a node is to be unfolded all of its parent nodes are also unfolded so that it is visible in the summary.

Shallowest unfold the shallowest available node first, choosing randomly if there is more than one. This would unfold node 1 in Figure 3 first, followed by either node 2 and 3 or node 3 and 2, etc.

Largest unfold the largest available node first, as measured by the number of tokens, breaking ties randomly. In Figure 3, this would unfold the class block (node 3) first, followed by the class Javadoc (node 4), etc.

Javadoc first unfold all Javadoc comments (in random order) and then fallback at random to an available node, unfolding methods last. This would unfold the Javadoc nodes 4, 6 in Figure 3 first and the method nodes 5, 8, 9 last (7 would already be unfolded as it is the parent of 10 and 11).

Each of the baselines represents a possible assumption that we can make about summarizing source code. Shallower nodes may be more important for the core meaning of the source code (*Shallowest* baseline) or that the largest nodes are more valuable in a summary (*Largest* baseline). The *Javadoc* baseline is representative of the de facto summarization method used in current IDEs such as Eclipse. Finally, we compare against our *TSSA* summarization method as formulated in Section 3.

Performance To assess the performance of our *TSSA* summarization system and baselines we randomly chose one of our two annotators and used their annotated test set (at both 75% and 50% compression ratios) as our gold standard for each file. This measures whether the output of the summarizer matches human judgements. For each file in the test set, we treated the folding problem as a binary classification, classifying each line of code as either unfolded (positive) or folded (negative). This enabled us to calculate the average accuracy, precision and recall of our summarizer across all test files and also F_1 as the harmonic mean of the average precision and recall. The resulting average performance metrics, at both 75% and 50% compression ratios, are given in Tables 3 and 4, respectively. As one can clearly see from the results, our proposed summarization system outperformed all the baselines, by a margin of about 15% when compared against the best performing baseline (Javadoc).

Moreover, our summarizer has very fast runtime, needing only a few seconds of CPU time to summarize an average file in the data set on a 2.66GHz Core 2 Quad machine.

Furthermore, in Table 5 we show the accuracy of the summarization methods averaged across the files for each test project. Once again, one can clearly see that for each project our TSSA summarizer significantly outperforms all the other

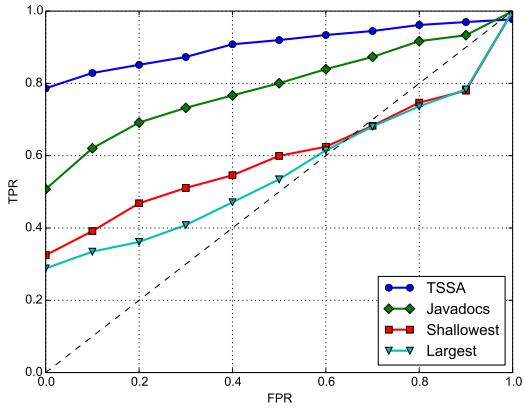

Figure 6: Averaged 11-point ROC curve for the summarizers at 0–100% compression ratios evaluated against the gold standard at 75%.

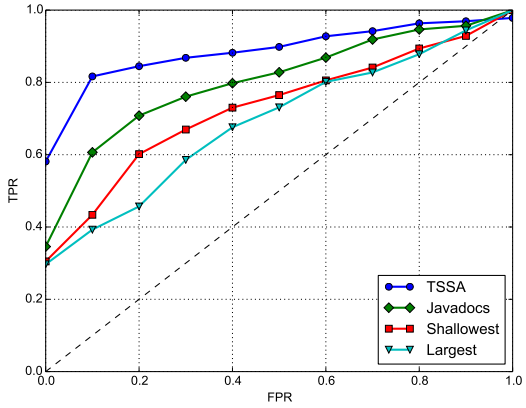

Figure 7: Averaged 11-point ROC curve for the summarizers at 0–100% compression ratios evaluated against the gold standard at 50%.

methods, showing that our proposed summarization method is not domain-dependent.

As a further test of whether the proposed summaries are plausible, we consider several classes of methods that are likely to be uninteresting, and verify that TSSA usually folds them. We show the percentage of times constructors, getters, setters and other generally uninteresting pattern-based method types were folded at 50% compression in Table 7. As one can see from the table, our summarizer folds these methods in most cases. Furthermore, when such methods are included in the summary, these exceptional methods turn out to be qualitatively more interesting, as we verify by manual examination. We show example snippets in Figures

project:	storm	elasticsearch	spring-framework	libgdx	bigbluebutton	netty
TSSA	**0.86**	**0.78**	**0.86**	**0.88**	**0.75**	**0.83**
Javadocs	0.70	0.70	0.71	0.70	0.58	0.76
Shallowest	0.57	0.53	0.49	0.51	0.56	0.62
Largest	0.73	0.43	0.49	0.54	0.52	0.54

Table 5: Comparison of per-line F_1 score across different open source projects, at a compression ratio of 50%. Averaged across project files, ordered best first.

```
public IndifferentAccessMap(IPersistentMap map) {
  setMap(map);
 }

IndexTemplateAlreadyExistsException(String name) {
  super("index_template [" + name
    + "] already exists");
  this.name = name;
 }
```

```
EsThreadPoolExecutor(int corePoolSize,
    int maximumPoolSize, long keepAliveTime,
    TimeUnit unit, BlockingQueue<Runnable> workQueue,
    ThreadFactory threadFactory) {
  this(corePoolSize, maximumPoolSize, keepAliveTime,
  unit, workQueue, threadFactory, new EsAbortPolicy());
 }

public PolygonMapObject(float[] vertices) {
  polygon = new Polygon(vertices);
 }
```

Figure 8: Snippets of constructors that were *unfolded* at 50% compression by our summarizer.

```
public Array<Rectangle> get(Rectangle r) {
  result.clear();
  int minX = max(0, cellX(r.x));
  int maxX = min(cols - 1, cellX(r.x + r.width));
  int minY = max(0, cellY(r.y));
  int maxY = min(rows - 1, cellY(r.y + r.height));
  for (int y = minY; y <= maxY; y++) {
   for (int x = minX; x <= maxX; x++) {
    int shv = getGridCell(x, y);
    result.addAll(grid.get(shv));
   }
  }
  return result;
 }

public <T> T getOption(ChannelOption<T> option) {
  if (option == SO_TIMEOUT) {
   return (T) Integer.valueOf(getSoTimeout());
  }
  return super.getOption(option);
 }

public void set(final ContactID c) {
  indexA = c.indexA;
  indexB = c.indexA;
  typeA = c.typeA;
  typeB = c.typeB;
 }

public <T> boolean
setOption(ChannelOption<T> option, T value) {
  validate(option, value;

  if (option == SO_TIMEOUT) {
   setSoTimeout((Integer) value);
  } else {
   return super.setOption(option, value);
  }
  return true;
 }
```

Figure 9: Snippets of getters and setters that were *unfolded* at 50% compression by our summarizer.

```
public static byte[]
readCompressedByteArray(DataInput in)
throws IOException {
  int length = in.readInt();
  if (length == -1)
   return null;
  byte[] buffer = new byte[length];
  in.readFully(buffer);
  GZIPInputStream gzi
    = new GZIPInputStream(new ByteArrayInputStream(
    buffer, 0, buffer.length));
  byte[] outbuf = new byte[length];
  ByteArrayOutputStream bos
    = new ByteArrayOutputStream();
  int len;
  while ((len = gzi.read(outbuf, 0,
    outbuf.length)) != -1) {
   bos.write(outbuf, 0, len);
  }
  byte[] decompressed = bos.toByteArray();
  bos.close();
  gzi.close();
  return decompressed;
 }

public void writeTo(StreamOutput out)
throws IOException {
  super.writeTo(out);
  if (out.getVersion().onOrBefore(Version.V_0_90_3)) {
   out.writeBoolean(false); // refresh flag
  }
  out.writeBoolean(full);
  out.writeBoolean(force);
 }
```

Figure 10: Snippets of readers and writers that were *unfolded* at 50% compression by our summarizer.

TSSA	Javadocs	Shallowest	Largest
0.89	0.81	0.73	0.69

Table 6: AUC for the ROC curve in Figure 7.

construct	get	set	put	is	read	write
74.6%	76%	92%	100%	100%	87.5%	70%

Table 7: The percentage of times specific method types in the test set were *folded* at 50% compression by our summarizer.

8–10. As one can see, the unfolded methods tend to exhibit unusual or non-standard behaviour in the method body, so much so that they cannot be easily summarized by their signature alone. This lends further evidence to the credibility and usefulness of our summarization approach.

It is also evident from the results that the Shallowest and Largest baselines perform poorly and it is not difficult to see why. The shallowest nodes tend to be code blocks delimiting methods and classes, which rarely contain core logic. Rather, the core logic tends to be nested in `if-else` statements, `for`/`while` loops and `try-catch` blocks. The Largest nodes may contain a substantial amount of code but rarely the core logic as the number of tokens is a bad indicator of code importance. To see this, consider header comments which contain many tokens (words) often stating the code copyright, or common class methods such as the equals method in Figure 1 which perform routine functions yet contain blocks with many identifiers. The Javadoc baseline, on the other hand, represents the de facto summarization method currently used in IDEs and therefore performs much better as one would expect.

As for statistical significance, we calculated two-tailed p-values using `sigf` [33]. We find that the difference between TSSA and the other methods is significant, obtaining values of $p < 0.01$ at both 75% and 50% compression ratios.

Finally, we compared the performance of the baselines and our summarization system at a range of compression ratios (effectively treating it as a threshold) against the gold standard at the fixed compression ratios of 75% and 50%. Figures 6 and 7 show the ROC curves with corresponding AUC values for Figure 7 in Table 6. Once again, we can see that our proposed summarization system substantially outperforms all the baselines with an AUC of 0.89 at 50% and a similar AUC at 75%.

6. CONCLUSIONS

We presented a novel fully unsupervised approach for extractive source code summarization formulating it as an optimal subtree problem on the source code's abstract syntax tree. Our proposed method builds on previous work using disjoint line- and term-level code summaries and introduces a new contiguous parse subtree as its summary. Moreover, in the process we develop a novel topic model for source code which has the potential for wider-ranging applications.

Our evaluation demonstrates that our proposed summarizer outperforms existing baselines, achieving an error reduction of 43% and even outperforming methods used as standard in modern IDEs, highlighting the immediate usefulness of our method.

In the future, we would like to extend our method to the summarization of multiple files. Furthermore, we are interested in integrating the term-based summaries developed by Haiduc et al. [15] with our own autofolding based summaries to obtain even better performance.

7. ACKNOWLEDGMENTS

This work was supported by Microsoft Research through its PhD Scholarship Programme and by the Engineering and Physical Sciences Research Council (grant number EP/K024043/1). We are also grateful to Rebecca Mason for allowing us to adapt her `TopicSum` implementation and would like to thank Brian Doll for useful discussions.

8. REFERENCES

[1] S. L. Abebe, S. Haiduc, P. Tonella, and A. Marcus. The effect of lexicon bad smells on concept location in source code. In *International Working Conference on Source Code Analysis and Manipulation (SCAM)*, pages 125–134. IEEE, 2011.

[2] M. Allamanis and C. Sutton. Mining source code repositories at massive scale using language modeling. In *International Workshop on Mining Software Repositories (MSR)*, pages 207–216. IEEE, 2013.

[3] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent Dirichlet allocation. *Journal of machine Learning research*, 3:993–1022, 2003.

[4] R. P. Buse and W. Weimer. Synthesizing API usage examples. In *International Conference on Software Engineering (ICSE)*, pages 782–792. IEEE, 2012.

[5] A. De Lucia, M. Di Penta, R. Oliveto, A. Panichella, and S. Panichella. Labeling source code with information retrieval methods: an empirical study. *Empirical Software Engineering*, pages 1–38, 2013.

[6] S. C. Deerwester, S. T. Dumais, T. K. Landauer, G. W. Furnas, and R. A. Harshman. Indexing by latent semantic analysis. *JASIS*, 41(6):391–407, 1990.

[7] B. Dit, M. Revelle, M. Gethers, and D. Poshyvanyk. Feature location in source code: a taxonomy and survey. *Journal of Software: Evolution and Process*, 25(1):53–95, 2013.

[8] Eclipse-Contributors. Eclipse JDT. `http://www.eclipse.org/jdt/`, 2013. Visited December 21, 2013.

[9] B. P. Eddy, J. A. Robinson, N. A. Kraft, and J. C. Carver. Evaluating source code summarization techniques: Replication and expansion. In *International Conference on Program Comprehension (ICPC)*, pages 13–22. IEEE, 2013.

[10] E. Erosheva, S. Fienberg, and J. Lafferty. Mixed-membership models of scientific publications. *Proceedings of the National Academy of Sciences of the United States of America*, 101(Suppl 1):5220–5227, 2004.

[11] M. Gethers, T. Savage, M. Di Penta, R. Oliveto, D. Poshyvanyk, and A. De Lucia. CodeTopics: which topic am I coding now? In *International Conference on Software Engineering (ICSE)*, pages 1034–1036. ACM, 2011.

[12] S. Guha, A. Moss, J. S. Naor, and B. Schieber. Efficient recovery from power outage. In *Symposium on Theory of Computing (STOC)*, pages 574–582. ACM, 1999.

[13] A. Haghighi and L. Vanderwende. Exploring content models for multi-document summarization. In *Association for Computational Linguistics (ACL)*, pages 362–370. ACL, 2009.

[14] S. Haiduc, J. Aponte, and A. Marcus. Supporting program comprehension with source code summarization. In *International Conference on Software Engineering (ICSE)*, volume 2, pages 223–226. IEEE, 2010.

[15] S. Haiduc, J. Aponte, L. Moreno, and A. Marcus. On the use of automated text summarization techniques for summarizing source code. In *Working Conference on Reverse Engineering (WCRE)*, pages 35–44, 2010.

[16] T. D. Hendrix, J. H. Cross II, L. A. Barowski, and K. S. Mathias. Visual support for incremental abstraction and refinement in Ada 95. In *ACM SIGAda Ada Letters*, volume 18, pages 142–147. ACM, 1998.

[17] A. Hindle, E. T. Barr, Z. Su, M. Gabel, and P. Devanbu. On the naturalness of software. In *International Conference on Software Engineering (ICSE)*, pages 837–847. IEEE, 2012.

[18] F. Jacob and R. Tairas. Code template inference using language models. In *Proceedings of the 48th Annual Southeast Regional Conference*, page 104. ACM, 2010.

[19] J. Kim, S. Lee, S.-W. Hwang, and S. Kim. Enriching documents with examples: A corpus mining approach. *Transactions on Information Systems (TOIS)*, 31(1):1, 2013.

[20] A. J. Ko, B. A. Myers, M. J. Coblenz, and H. H. Aung. An exploratory study of how developers seek, relate, and collect relevant information during software maintenance tasks. *IEEE Transactions on Software Engineering*, 32(12):971–987, 2006.

[21] B. Kullbach and V. Riediger. Folding: An approach to enable program understanding of preprocessed languages. In *Working Conference on Reverse Engineering (WCRE)*, pages 3–12. IEEE, 2001.

[22] T. D. LaToza, G. Venolia, and R. DeLine. Maintaining mental models: a study of developer work habits. In *International Conference on Software Engineering (ICSE)*, pages 492–501. ACM, 2006.

[23] D. Lawrie, C. Morrell, H. Feild, and D. Binkley. What's in a name? A study of identifiers. In *International Conference on Program Comprehension (ICPC)*, pages 3–12. IEEE, 2006.

[24] C.-Y. Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, pages 74–81, 2004.

[25] S. Mani, R. Catherine, V. S. Sinha, and A. Dubey. Ausum: approach for unsupervised bug report summarization. In *Foundations of Software Engineering (FSE)*, page 11. ACM, 2012.

[26] A. Marcus and J. I. Maletic. Recovering documentation-to-source-code traceability links using latent semantic indexing. In *International Conference on Software Engineering (ICSE)*, pages 125–135, Washington, DC, USA, 2003. IEEE.

[27] A. Marcus, A. Sergeyev, V. Rajlich, and J. I. Maletic. An information retrieval approach to concept location in source code. In *Working Conference on Reverse Engineering (WCRE)*, pages 214–223. IEEE, 2004.

[28] D. Mimno, W. Li, and A. McCallum. Mixtures of hierarchical topics with pachinko allocation. In *International Conference on Machine Learning (ICML)*, pages 633–640. ACM, 2007.

[29] L. Moreno, J. Aponte, G. Sridhara, A. Marcus, L. Pollock, and K. Vijay-Shanker. Automatic generation of natural language summaries for Java classes. In *International Conference on Program Comprehension (ICPC)*, pages 23–32. IEEE, 2013.

[30] D. Movshovitz-Attias and W. W. Cohen. Natural language models for predicting programming comments. In *Association for Computational Linguistics (ACL)*, 2013.

[31] A. T. Nguyen, T. T. Nguyen, and T. N. Nguyen. Lexical statistical machine translation for language migration. In *Joint Meeting on Foundations of Software Engineering (FSE)*. ACM, 2013.

[32] T. T. Nguyen, A. T. Nguyen, H. A. Nguyen, and T. N. Nguyen. A statistical semantic language model for source code. In *Joint Meeting on Foundations of Software Engineering (FSE)*, pages 532–542. ACM, 2013.

[33] S. Padó. *User's guide to `sigf`: Significance testing by approximate randomisation*, 2006.

[34] S. Rastkar, G. C. Murphy, and G. Murray. Summarizing software artifacts: a case study of bug reports. In *International Conference on Software Engineering (ICSE)*, pages 505–514. ACM, 2010.

[35] S. Rugaber, N. Chainani, O. Nnadi, and K. Stirewalt. A conceptual model for folding. Technical Report GT-CS-08-09, Georgia Institute of Technology, 2008.

[36] T. Savage, B. Dit, M. Gethers, and D. Poshyvanyk. Topic XP: Exploring topics in source code using latent Dirichlet allocation. In *International Conference on Software Maintenance (ICSM)*, pages 1–6. IEEE, 2010.

[37] K. Spärck Jones. Automatic summarising: The state of the art. *Information Processing & Management*, 43(6):1449–1481, 2007.

[38] G. Sridhara, E. Hill, D. Muppaneni, L. Pollock, and K. Vijay-Shanker. Towards automatically generating summary comments for Java methods. In *International Conference on Automated Software Engineering (ASE)*, pages 43–52. ACM, 2010.

[39] G. Sridhara, L. Pollock, and K. Vijay-Shanker. Automatically detecting and describing high level actions within methods. In *International Conference on Software Engineering (ICSE)*, pages 101–110, 2011.

[40] J. Starke, C. Luce, and J. Sillito. Searching and skimming: An exploratory study. In *International Conference on Software Maintenance (ICSM)*, pages 157–166. IEEE, 2009.

[41] M. Steyvers and T. Griffiths. Probabilistic topic models. *Handbook of latent semantic analysis*, 427(7):424–440, 2007.

[42] M.-A. Storey. Theories, methods and tools in program comprehension: Past, present and future. In *International Workshop on Program Comprehension (IWPC)*, pages 181–191. IEEE, 2005.

[43] S. W. Thomas. Mining software repositories using topic models. In *International Conference on Software*

Engineering (ICSE), pages 1138–1139. ACM, 2011.

[44] H. M. Wallach. *Structured topic models for language*. PhD thesis, 2008.

[45] J. Wang, Y. Dang, H. Zhang, K. Chen, T. Xie, and D. Zhang. Mining succinct and high-coverage API usage patterns from source code. In *International Workshop on Mining Software Repositories (MSR)*, pages 319–328. IEEE Press, 2013.

[46] T. Xie and J. Pei. MAPO: Mining API usages from open source repositories. In *International Workshop on Mining Software Repositories (MSR)*, pages 54–57. ACM, 2006.

[47] A. T. T. Ying and M. P. Robillard. Code fragment summarization. In *Joint Meeting on Foundations of Software Engineering (FSE)*, pages 655–658. ACM, 2013.